On the Hardness of the Finite Field Isomorphism Problem

Dipayan Das[0000−0002−7681−1868] and Antoine Joux[0000−0003−2682−6508]

CISPA Helmholtz Center for Information Security, Saarbrücken, Germany.
dipayan.das@cispa.de, joux@cispa.de

Abstract. The finite field isomorphism (FFI) problem was introduced in PKC'18, as an alternative to average-case lattice problems (like LWE, SIS, or NTRU). As an application, the same paper used the FFI problem to construct a fully homomorphic encryption scheme. In this work, we prove that the decision variant of the FFI problem can be solved in polynomial time for any field characteristics $q = \Omega(\beta n^2)$, where q, β, n parametrize the FFI problem. Then we use our result from the FFI distinguisher to propose polynomial-time attacks on the semantic security of the fully homomorphic encryption scheme. Furthermore, for completeness, we also study the search variant of the FFI problem and show how to state it as a q-ary lattice problem, which was previously unknown. As a result, we can solve the search problem for some previously intractable parameters using a simple lattice reduction approach.

1 Introduction

The Finite Field Isomorphism (FFI) problem has been introduced in [4] as a new hard problem to study post-quantum cryptography. Informally, it states the following.

For a hidden element $\boldsymbol{x}$ (with sparse minimal polynomial) in the finite field $\mathbb{F}_{q^n}$, if small β-bounded linear combinations of powers of $\boldsymbol{x}$ are given, in terms of powers of a uniform generator $\boldsymbol{y}$, it is hard to recover $\boldsymbol{x}$.

The decisional version of the problem states the following.

Given the $\boldsymbol{y}$-basis representation of finite field elements, it is hard to decide whether they are picked from the FFI *distribution or the uniform distribution, with a non-negligible advantage over random guessing.*

The FFI assumption is based on the fact that the basis transformation converts "good" representations to "bad" representations of $\mathbb{F}_{q^n}$. At a high level of abstraction, the heuristics of the FFI problem is comparable to many lattice problems, which involve recovering a "good" secret basis from a "bad" public basis (example, [7,8]). However, despite this high-level similarity, the details are quite different and a dedicated security analysis is required.

In the papers [4,9], the authors thoroughly analyzed the generic hardness of the FFI problem. From their analysis, the best known attack for the decisional problem has $2^{O(n)}$ time complexity, whereas the best known attack for the search problem has $O(n!)$ time complexity. Based on their analysis, they proceed to

propose a fully homomorphic encryption scheme [4] and a signature scheme [9] as applications of the FFI problem.

1.1 Our Contribution

This paper re-examines the hardness of the FFI problem in both its decisional and computational versions. We use basic finite field theory to study the hardness of the FFI problem.

In Section 4, we prove the values of the trace of the hidden polynomial $\boldsymbol{x}$-basis are bounded in absolute value by n. The proof is based on combinatorial techniques. Thus, by a linearity argument, the trace of FFI samples can be bounded in absolute value by βn^2. This observation provides a polynomial-time distinguisher to solve the decisional version of the FFI problem for any field with characteristic $q \geq 4\beta n^2$.

In Section 5, we complement the attack on the decisional FFI problem by breaking the semantic security of the fully homomorphic encryption scheme from [4]. First, we give a simple semantic attack by using the bound on the trace of the ciphertext. Then we provide an additional semantic attack that offers a slight improvement on the lower bound on the field characteristic q, using the distribution of the trace of the inverse of the finite field elements. Both attacks apply to the proposed parameter of the scheme.

In Section 6, we exploit the notion of dual basis of a finite field basis to solve the computational FFI problem. By definition, the trace of any basis element with respect to its dual basis can be expressed as the Kronecker delta function. Consequently, the traces of FFI samples with respect to the dual $\boldsymbol{x}$-basis are bounded by β. Using this observation, we recover the dual $\boldsymbol{x}$-basis from the given FFI samples $\boldsymbol{A}_i$. This is done by reducing an adequate lattice built from traces of well-chosen finite field elements. We also show that a partial recovery of the dual $\boldsymbol{x}$-basis can be leverage into a full cryptanalysis with good probability.

Finally, we provide some experiments on lattice reduction to find the shortest vectors in this lattice.

2 Preliminaries

2.1 Notations

The parameter q denotes a (moderately large) prime integer throughout the paper.

The finite field with q elements is denoted by $\mathbb{F}_q$. All vectors are in columns and are denoted with bold letters. We identify polynomials and vectors as being the same data type using the coefficient embedding. For any vector $\boldsymbol{v}$, we write $\|\boldsymbol{v}\|$ for the ℓ_∞ norm of $\boldsymbol{v}$, and $\|\boldsymbol{v}\|_2$ for the ℓ_2 norm of $\boldsymbol{v}$. Representatives of the elements of $\mathbb{F}_q$ are centered around zero, i.e. chosen in the interval $\left[-\frac{q-1}{2}, \frac{q-1}{2}\right]$. The rationale for using this representation is that it is much better adapted to the goal of obtaining vectors with short norms.

2.2 Reminders from Finite Field Theory

For every prime q and every positive integer n, there exists a unique finite field with q^n elements. It is denoted by $\mathbb{F}_{q^n}$. The prime q and the integer n are respectively called the characteristic and degree of the finite field.

Let $\boldsymbol{f}(\boldsymbol{x})$ and $\boldsymbol{F}(\boldsymbol{y})$ be two irreducible polynomials of degree n over $\mathbb{F}_q$. We can construct two isomorphic representations of the finite field $\mathbb{F}_{q^n}$ as $\mathbb{X} := \mathbb{F}_q[\boldsymbol{x}]/(\boldsymbol{f}(\boldsymbol{x}))$ and $\mathbb{Y} := \mathbb{F}_q[\boldsymbol{y}]/(\boldsymbol{F}(\boldsymbol{y}))$. Every element of $\mathbb{F}_{q^n}$ can be uniquely represented by a polynomial in $\boldsymbol{x}$ with coefficients in $\mathbb{F}_q$ and degree less than n. Similarly, there is a representation in terms of $\boldsymbol{y}$. In other words, the set

$$\{\mathbf{1}, \boldsymbol{x}, \ldots, \boldsymbol{x}^{n-1}\}$$

is a basis of $\mathbb{F}_{q^n}$, viewed as a vector space over $\mathbb{F}_q$. This basis is called the $\boldsymbol{x}$-polynomial-basis or $\boldsymbol{x}$-basis for short. For ease of reading, we denote finite field elements known in the $\boldsymbol{x}$-basis by small letters and elements known in the $\boldsymbol{y}$-basis by capital letters.

To explicit an isomorphism between these two representations of $\mathbb{F}_{q^n}$, it suffices to know the representation in the $\boldsymbol{x}$-basis of a root of $\boldsymbol{F}(\boldsymbol{y})$ or conversely the $\boldsymbol{y}$-representation of a root of $\boldsymbol{f}(\boldsymbol{x})$. Note that each of the two polynomials has n distinct roots, which are images of each other by the q-th power Frobenius map.

For every element $\boldsymbol{\alpha} \in \mathbb{F}_{q^n}$, its conjugates are obtained by repeatedly applying the Frobenius map, i.e. they are $\boldsymbol{\alpha}$, $\boldsymbol{\alpha}^q$, $\ldots$, $\boldsymbol{\alpha}^{q^{(n-1)}}$. They are distinct if and only if the minimal polynomial of $\boldsymbol{\alpha}$ has degree n. The trace of an element in $\mathbb{F}_{q^n}$ is defined as the sum of all its conjugates:

$$\mathsf{Tr}(\boldsymbol{\alpha}) := \boldsymbol{\alpha} + \boldsymbol{\alpha}^q + \cdots + \boldsymbol{\alpha}^{q^{(n-1)}} \in \mathbb{F}_q$$

The trace function is linear, i.e.

$$\mathsf{Tr}(\boldsymbol{\alpha} + c\boldsymbol{\beta}) = \mathsf{Tr}(\boldsymbol{\alpha}) + c\mathsf{Tr}(\boldsymbol{\beta})$$

for any $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{F}_{q^n}$ and $c \in \mathbb{F}_q$.

Moreover, for every linear map L from $\mathbb{F}_{q^n}$ to $\mathbb{F}_q$, there exists a unique element $\boldsymbol{\beta}$ in $\mathbb{F}_{q^n}$ such that:

$$\forall \boldsymbol{\alpha} \in \mathbb{F}_{q^n} : L(\boldsymbol{\alpha}) = \mathsf{Tr}(\boldsymbol{\beta} \cdot \boldsymbol{\alpha}).$$

We denote this linear function by $L_{\boldsymbol{\beta}}$.

To every basis $\boldsymbol{\omega}_1, \boldsymbol{\omega}_2, \ldots, \boldsymbol{\omega}_n$ of $\mathbb{F}_{q^n}$ we associate a dual basis [1] $\widehat{\boldsymbol{\omega}}_1, \widehat{\boldsymbol{\omega}}_2, \ldots, \widehat{\boldsymbol{\omega}}_n \in \mathbb{F}_{q^n}$, defined as the unique one which satisfies:

$$\mathsf{Tr}(\boldsymbol{\omega}_i \widehat{\boldsymbol{\omega}}_j) = \delta_i^j$$

where δ_i^j is the *Kronecker delta* function. From this definition, it is clear that the bidual of a basis, i.e. the dual of its dual, is the basis itself.

[1] Note that this notion of the dual basis of a finite field does not correspond to the idea of the dual basis of a lattice. This paper only uses the term dual basis to refer to the former notion.

2.3 Lattice Reduction

Given a (full rank) matrix $\boldsymbol{B} \in \mathbb{Z}^{d \times d}$, the lattice $\mathcal{L}$ generated by the basis $\boldsymbol{B}$ is the set $\mathcal{L}(\boldsymbol{B}) := \{\boldsymbol{Bz} : \boldsymbol{z} \in \mathbb{Z}^d\}$, d is the lattice dimension. A lattice is called q-ary if it contains $q\mathbb{Z}^d$ as a sublattice. The volume of a lattice $\mathcal{L}(\boldsymbol{B})$ is defined as $\mathbf{Vol}(\mathcal{L}) := |\det(\boldsymbol{B})|$. Any lattice of dimension $d \geq 2$ has infinitely many bases that generate the same lattice, and any two bases $\boldsymbol{B}, \boldsymbol{B'}$ are related by a unimodular matrix $\boldsymbol{U}$ such that $\boldsymbol{B} = \boldsymbol{B'U}$. Note that the unimodular matrix stands on the right because we use the convention of having vectors in columns. The volume of a lattice is independent of the choice of lattice basis.

For a random lattice $\mathcal{L}$, the Gaussian heuristic estimates the Euclidean norm of the shortest non-zero vector in the lattice, which is approximately $\sqrt{\frac{d}{2\pi e}}\mathbf{Vol}(\mathcal{L})^{1/d}$ [5].

For any basis $\boldsymbol{B}$, we write $\boldsymbol{B}^*$ to represent the Gram-Schmidt orthogonalization (GSO) of $\boldsymbol{B}$, where the i-th vector of $\boldsymbol{B}^*$ is given by $\boldsymbol{b}_i^* := \pi_i(\boldsymbol{b}_i)$. Here, the notation π_i denotes the projection of a vector orthogonally to the vector subspace spanned by $\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_{i-1}$.

A central problem in the algorithmics of lattices is to find shortest non-zero vectors (SVP) from a lattice basis $\boldsymbol{B}$. This can be a handy tool in cryptanalysis. The most widely used lattice reduction algorithm is LLL [11] which is polynomial-time but only yields an approximation of SVP within an exponential factor. Since this can be insufficient for cryptanalysis, it is standard practice to use slower algorithms that produce better approximations.

For our needs, we use the implementation of the blockwise Korkine-Zolotarev (BKZ) algorithm provided with the fplll software [13].

2.4 Semantic Attack of an encryption scheme

An encryption scheme $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ **that only encrypts bits** (i.e. with message space $\{0,1\}$) has (t, δ) attack against the semantic security if there exists an adversary $\mathcal{A}$ winning the following game against a challenger $\mathcal{C}$.

- $\mathcal{C}$ samples $m \hookleftarrow \{0,1\}$, $(\mathsf{pk}, \mathsf{sk}) \hookleftarrow \mathsf{KeyGen}(1^\lambda)$.
- $\mathcal{C}$ gives pk, $c := \mathsf{Enc}(\mathsf{pk}, m)$ to $\mathcal{A}$.
- $\mathcal{A}$ outputs $m' \in \{0,1\}$.

$\mathcal{A}$ wins the game if $\mathcal{A}$ has running time t and advantage δ, where the advantage is defined by

$$|\Pr[m' = m] - \Pr[m' \neq m]|$$

3 Finite Field Isomorphism Problem

This section formally describes the FFI problem in both its computational and decisional forms.

Let $\mathbb{X}, \mathbb{Y}$ be two representations of the finite field $\mathbb{F}_{q^n}$ as before. In the rest of the paper, we assume $n \geq 50$ to be out of range of easy exhaustive search attacks. The defining polynomial of $\mathbb{X}$ is sampled uniformly from the set of all sparse irreducible polynomials of the form $\boldsymbol{x}^n + \boldsymbol{g}(\boldsymbol{x})$ with $\deg \boldsymbol{g}(\boldsymbol{x}) \leq \lfloor n/2 \rfloor$ and $\|\boldsymbol{g}(\boldsymbol{x})\| \leq 1$, i.e. $\boldsymbol{g}(\boldsymbol{x})$ has ternary coefficients. The defining polynomial of $\mathbb{Y}$ is sampled uniformly from the set of arbitrary monic irreducible polynomials. Let $\boldsymbol{\phi}(\boldsymbol{y})$ be an isomorphism from $\mathbb{X}$ to $\mathbb{Y}$. Note that there is an efficient algorithm to compute an isomorphism between the two representations [2].[2] Let χ_β be a distribution over $\mathbb{X}$ that samples polynomials $\boldsymbol{a}_i(\boldsymbol{x})$ with $\|\boldsymbol{a}_i(\boldsymbol{x})\| \leq \beta$. Let $\boldsymbol{A}_i(\boldsymbol{y})$ be the corresponding image of $\boldsymbol{a}_i(\boldsymbol{x})$ under the isomorphism $\boldsymbol{\phi}$.

Definition 1 (Computational Finite Field Isomorphism Problem ($\mathsf{CFFI}_{q,k,n,\beta}$)**).** *Given $\mathbb{Y}$ by $\boldsymbol{F}(\boldsymbol{y})$ and k samples $\boldsymbol{A}_1(\boldsymbol{y}), \ldots, \boldsymbol{A}_k(\boldsymbol{y})$ recover $\boldsymbol{f}(\boldsymbol{x})$.*

Definition 2 (Decisional Finite Field Isomorphism Problem ($\mathsf{DFFI}_{q,k,n,\beta}$)**).** *Given $\mathbb{Y}$ by $\boldsymbol{F}(\boldsymbol{y})$ and k samples $\boldsymbol{B}_1(\boldsymbol{y}), \boldsymbol{B}_2(\boldsymbol{y}), \ldots, \boldsymbol{B}_k(\boldsymbol{y})$ that are either sampled from* FFI *distribution (having pre-images bounded by β) or sampled uniformly at random in $\mathbb{Y}$, $\mathsf{DFFI}_{q,k,n,\beta}$ problem is to distinguish, with some non-negligible advantage, the correct source distribution of the samples $\boldsymbol{B}_i(\boldsymbol{y})$.*

Note that the sparsity constraint on the defining polynomial of $\mathbb{X}$ is not directly included in the definition of the FFI problem given in [4,9]. However, the noise growth analysis of [4, Appendix B] explicitly rewrites[3] $\boldsymbol{f}(\boldsymbol{x})$ as $\boldsymbol{x}^n + \boldsymbol{f}'(\boldsymbol{x})$ and proceeds to bound the noise-growth during multiplication in $\mathbb{X}$ under the assumption that the degree d of $\boldsymbol{f}'$ satisfies $d < n/2$. For clarity, we instead chose to directly include this low-degree constraint as part of the definition.

3.1 Previous Attacks

In this section, we briefly describe all the attacks that have been considered for both decisional and computational FFI problem in [4,9].

Decisional Finite Field Isomorphism Problem

Lattice Attack

The decisional FFI problem could be solved by predicting if there is any good representation of the given samples, which is very unlikely for uniform samples. To achieve this, the authors of [4,9] suggested lattice reduction on the q-ary lattice

$$\mathcal{L}_{\boldsymbol{A},q} := \{\boldsymbol{a}_i \in \mathbb{Z}^k : \boldsymbol{A}\boldsymbol{\Psi}_i = \boldsymbol{a}_i \bmod q \text{ for some } \boldsymbol{\Psi}_i \in \mathbb{Z}^n\}$$

[2] In practice, SageMath provides the `FiniteFieldHomomorphism_generic(Hom(.))` function available under `sage.rings.finite_rings.hom_finite_field` package for this task.

[3] In the rewriting, $\boldsymbol{f}'(\boldsymbol{x})$ does not denote the derivative of $\boldsymbol{f}$ but an auxiliary polynomial. In our definition, we use the notation $\boldsymbol{g}(\boldsymbol{x})$ to avoid any possible confusion with the derivative.

with each row of the matrix $\boldsymbol{A}$ generated from the given $\boldsymbol{y}$-basis representations of samples. For FFI samples, there are unusually short vectors in the lattice that corresponds to the $\boldsymbol{x}$-basis (or small linear combination of $\boldsymbol{x}$-basis) representation. For uniform samples, it is highly unlikely to have such short vectors in the lattice.

Computational Finite Field Isomorphism Problem

Hybrid attack
The authors of [4,9] propose to find the shortest vectors in the lattice $\mathcal{L}_{\boldsymbol{A},q}$, in the hope that they correspond to the coefficients of some powers of $\boldsymbol{x}$. Since the shortest vectors appear in a somewhat random-looking order, the authors suggested adding a combinatorial algorithm to resolve the ordering issue and recover the $\boldsymbol{x}$-basis representations of the FFI samples. This gives an attack to the computational FFI problem. They estimate the cost of the combinatorial step to be $O(n!)$, thus infeasible. One might argue that this could possibly be improved by some form of meet-in-the-middle to $O(\sqrt{n!})$. We do not examine this direction since we show in Section 6 that we can get rid of the combinatorial step altogether.

Non-linear attack
The non-linear attack involves solving the non-linear system of equations to recover the hidden isomorphism $\boldsymbol{\phi}$ using Gröbner basis computation. An adversary can solve for $2n-2$ unknowns of $(\boldsymbol{\phi}, (\boldsymbol{a}_i(\boldsymbol{x}))$ from the equation $\boldsymbol{\phi}(\boldsymbol{a}_i(\boldsymbol{x})) = \boldsymbol{A}_i(\boldsymbol{y})$. Solving such an equation is believed to be hard.

4 Proposed Attack on the Decisional FFI problem

This section proposes a new polynomial-time attack on the DFFI problem. We show that when a sample $\boldsymbol{A}_i(\boldsymbol{y})$ comes from the FFI distribution, the underlying trace of $\boldsymbol{A}_i(\boldsymbol{y})$ is bounded by a small multiple of n^2. We use this fact to mount a distinguishing attack on the DFFI problem.

Lemma 1. *Let* $\boldsymbol{f}(\boldsymbol{x}) := \boldsymbol{x}^n + \sigma_1\boldsymbol{x}^{n-1} + \sigma_2\boldsymbol{x}^{n-2} + \cdots + \sigma_n$ *be the defining polynomial of* $\mathbb{X}$*, where* $\sigma_i \in \{-1, 0, 1\}$ *for* $\lceil n/2 \rceil \leq i \leq n$*,* 0 *otherwise*[4]*.*

Then

$$\mathsf{Tr}(\boldsymbol{x}^i) \equiv \begin{cases} n \pmod q, & \text{if } i = 0 \\ 0 \pmod q, & \text{if } 1 \leq i \leq \lceil n/2 \rceil - 1 \\ \pm i \pmod q, & \text{if } \sigma_i \neq 0 \text{ and } \lceil n/2 \rceil \leq i \leq n-1 \\ 0 \pmod q, & \text{if } \sigma_i = 0 \text{ and } \lceil n/2 \rceil \leq i \leq n-1 \end{cases}$$

Proof. The definition of trace function gives $\mathsf{Tr}(\boldsymbol{x}^i) = n$ when $i = 0$. To prove the Lemma for $i > 0$, let us first recall the Girard-Newton identities relating the sum

[4] Indeed, the smallest i with $\sigma_i \neq 0$ satisfies $i + \deg \boldsymbol{g} = n$. Since $\deg \boldsymbol{g} = \lfloor n/2 \rfloor$ this corresponds to $i = \lceil n/2 \rceil$.

of powers to symmetric polynomials. Given n arbitrary numbers $w_0, \ldots, w_{n-1}$ in an arbitrary ring, define their symmetric polynomials as usual by: $\sigma_1 = \sum_{i=0}^{n-1} w_i$, $\sigma_2 = \sum_{0 \le i < j < n} w_i w_j$, $\sigma_3 = \sum_{0 \le i < j < k < n} w_i w_j w_k$, $\ldots$, $\sigma_n = \prod_{i=0}^{n-1} w_i$. Then, for any $1 \le d < n$, we have:

$$\sum_{i=0}^{n-1} w_i^d = (-1)^d \sum_{\substack{r_1 + 2r_2 + \cdots + dr_d = d \\ r_i \in \mathbb{N}}} d \cdot \frac{(r_1 + r_2 + \cdots + r_d - 1)!}{r_1! r_2! \ldots r_d!} \prod_{j=1}^{d} (-\sigma_j)^{r_j} \quad (1)$$

Note that the coefficients in this equation, while written as fractions for notational purposes are, in fact, integers. As a consequence, the identity holds in every ring. In particular, when working modulo q as we are. However, to avoid any potential division by 0, the coefficients should first be computed as exact integers and only reduced modulo q afterwards.

The set of all the roots of $\boldsymbol{f}(\boldsymbol{x})$ are given by

$$\{\boldsymbol{\alpha}_0 := \boldsymbol{x}, \boldsymbol{\alpha}_1 := \boldsymbol{x}^q, \ldots, \boldsymbol{\alpha}_{n-1} := \boldsymbol{x}^{q^{n-1}}\}$$

Let now the $(\sigma_j)_{1 \le j \le n}$ denote the n symmetric polynomials in these roots. We know that we can write

$$\boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{x}^n - \sigma_1 \boldsymbol{x}^{n-1} + \sigma_2 \boldsymbol{x}^{n-2} - \ldots (-1)^n \sigma_n.$$

Thus, $\sigma_1 = \sigma_2 = \cdots = \sigma_{\lceil n/2 \rceil - 1} = 0$.

Since, $\mathsf{Tr}(\boldsymbol{x}^i)$ is a sum of i-th power of $\boldsymbol{\alpha}_j$s, we can use the above identity to express it in terms of the σ_js. Depending on the value of i, two cases arise:

Case 1 ($1 \le i \le \lceil n/2 \rceil - 1$). Since all contributions include a σ_j for $j \le i$ with value zero, there is no non-zero term in the sum of the right-hand side of Equation (1), we have

$$\mathsf{Tr}(\boldsymbol{x}^i) = 0$$

Case 2 ($\lceil n/2 \rceil \le i \le n - 1$). Again, since $\sigma_j = 0$ for $1 \le j \le \lceil n/2 \rceil - 1$, there is exactly one element in the set $\{(r_1, r_2, \ldots, r_i) : \sum_{l=1}^{i} l r_l = i\}$ that contributes in the sum of the right-hand side of Equation (1), namely $(r_1 = 0, r_2 = 0, \ldots, r_i = 1)$. Indeed, the sum of two contributions above $\lceil n/2 \rceil$ is always greater than i. This gives

$$\mathsf{Tr}(\boldsymbol{x}^i) = \begin{cases} \pm i \text{ for non-zero } \sigma_i, \\ 0 \text{ for } \sigma_i = 0 \end{cases}$$

Lemma 2. *Let $\boldsymbol{A}_i(\boldsymbol{y})$ be an $\mathsf{FFI}_{q,k,n,\beta}$ sample. Then $|\mathsf{Tr}(\boldsymbol{A}_i(\boldsymbol{y}))| \le 0.39\beta n^2$.*

Proof. Let $\boldsymbol{a}_i(\boldsymbol{x})$ be the representation of $\boldsymbol{A}_i(\boldsymbol{y})$ in the $\boldsymbol{x}$-basis. Since the trace of a finite field element is invariant to the basis representation, both $\boldsymbol{a}_i(\boldsymbol{x})$ and $\boldsymbol{A}_i(\boldsymbol{y})$ must have the same trace. So in order to bound the trace of $\boldsymbol{A}_i(\boldsymbol{y})$, it is

sufficient to bound the trace of $\boldsymbol{a}_i(\boldsymbol{x})$. By the linearity of the trace and by the previous Lemma, we have

$$\begin{aligned}|\mathsf{Tr}(\boldsymbol{A}_i(\boldsymbol{y}))| &\leq \beta \sum_{j=\lceil n/2 \rceil}^{n} j \\ &\leq 0.39\beta n^2 \text{ for } n \geq 50.\end{aligned}$$

Theorem 1. *Let $q \geq 1.56\beta n^2$, then there exists a polynomial-time algorithm with advantage $1 - \Omega(1/2^k)$ to distinguish the $\mathsf{DFFI}_{q,k,n,\beta}$ problem.*

Proof. Let $\boldsymbol{B}_1(\boldsymbol{y}), \boldsymbol{B}_2(\boldsymbol{y}), \ldots, \boldsymbol{B}_k(\boldsymbol{y})$ be the given k samples to distinguish between FFI distribution and uniform distribution. The distinguisher finds the correct distribution of the samples by computing the trace of the samples. In a finite field, the trace function is uniformly distributed. Thus for a uniform sample $\boldsymbol{B}_i(\boldsymbol{y})$, $\mathsf{Tr}(\boldsymbol{B}_i(\boldsymbol{y}))$ is uniformly distributed over $\mathbb{F}_q$. For an FFI sample $\boldsymbol{B}_i(\boldsymbol{y})$, by the previous Lemma, $|\mathsf{Tr}(\boldsymbol{B}_i(\boldsymbol{y}))| \leq 0.39\beta n^2$. Combining the number of samples and condition on q, the distinguisher outputs 1 when the samples come from FFI distribution with probability 1, and outputs 1 when the samples come from uniform distribution with probability at most $1/2^k$.

It is only left to show that the distinguisher is indeed polynomial-time. The running time of the attack is dominated by trace computation of finite field elements. Since the trace of a finite field element can be computed efficiently in time $n^{1+o(1)} \log^{2+o(1)} q$ using iterated Frobenius [10,12], this is polynomial-time.

5 Proposed Semantic Attack on the fully homomorphic encryption scheme

In this section, we propose a polynomial-time attack on the semantic security of the fully homomorphic encryption scheme $\mathcal{E} := (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ from [4]. The working principle of the scheme is given below.

- $\mathsf{KeyGen}(1^\lambda)$: Generate the FFI parameters $\Xi := (n, q, \beta)$ as a function of λ, and two representations of the finite field by sampling $\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{F}(\boldsymbol{y})$ with an isomorphism $\boldsymbol{\phi}$ like before. Choose two integers (S, s) satisfying $\binom{S}{s} \geq 2^\lambda$. Sample S many $\boldsymbol{c}_i(\boldsymbol{x})$ from the distribution χ_β and construct $\boldsymbol{C}_i(\boldsymbol{y}) := p\boldsymbol{\phi}(\boldsymbol{c}_i(\boldsymbol{x}))$ for fixed constant $p := 2$. In the rest of the section, we assume p is equal to 2 as in [4].
  The secret key is $\mathsf{sk} := (\Xi, \boldsymbol{\phi}, \boldsymbol{f}(\boldsymbol{x}))$.
  The public key is $\mathsf{pk} := (\Xi, \boldsymbol{C}_1(\boldsymbol{y}), \boldsymbol{C}_2(\boldsymbol{y}), \ldots, \boldsymbol{C}_S(\boldsymbol{y}), \boldsymbol{F}(\boldsymbol{y}), s, p)$.

- $\mathsf{Enc}(m, \mathsf{pk})$: The encryption of a message $m \in \{0, 1\}$ is

$$\boldsymbol{C} := \sum_{i \in [s]} \boldsymbol{C}_i(\boldsymbol{y}) + m$$

  for uniformly random s samples $\boldsymbol{C}_i(\boldsymbol{y})$.

– $\mathsf{Dec}(\boldsymbol{C}, \mathsf{sk})$: The decryption recovers m by computing

$$m' := p \sum_{i \in [s]} \boldsymbol{c}_i(\boldsymbol{x}) + m \bmod p$$

using the inverse of the secret isomorphism $\boldsymbol{\phi}$.

– $\mathsf{Eval}(C, \boldsymbol{C}^{(1)}, \boldsymbol{C}^{(2)}, \ldots, \boldsymbol{C}^{(l)})$: The homomorphic evaluation of ciphertexts of a circuit C with gates $(+, \times)$ are done using homomorphic addition and multiplication (with noise management) on $\boldsymbol{C}^{(i)}$s. It is shown that for $q = 2^{n^\epsilon}$ with $\epsilon \in (0, 1)$, the above encryption scheme $\mathcal{E}$ is fully homomorphic using circular security and bootstrapping techniques (Theorem 3 of [4]).

The result of Theorem 1 invalidates the semantic security of the fully homomorphic encryption scheme $\mathcal{E}$ (Theorem 1 of [4]). The next Theorem gives a polynomial-time algorithm to break the semantic security of $\mathcal{E}$.[5]

Theorem 2. *Let $q > 0.44s\beta^2n^5 + 0.78\beta n^2$, then there exists a deterministic polynomial time attack against the semantic security of the fully homomorphic encryption scheme $\mathcal{E}$ defined as above.*

Proof. Let the challenger $\mathcal{C}$ give the public key $\mathsf{pk} \hookleftarrow \mathsf{KeyGen}(1^\lambda)$ and an encryption $\boldsymbol{C}$ of a message $m \hookleftarrow \{0, 1\}$ to the adversary $\mathcal{A}$. Then

$$\boldsymbol{C} = p \sum_{i \in [s]} \boldsymbol{\phi}(\boldsymbol{c}_i(\boldsymbol{x})) + m \tag{2}$$

$\mathcal{A}$ wins the semantic game by the following analysis.

We consider the following two cases for the choice of n.

1. When $p = 2$ is not a divisor of n.

   Note that $\sum_{i \in [s]} \boldsymbol{\phi}(\boldsymbol{c}_i(\boldsymbol{x}))$ is an $\mathsf{FFI}_{q,1,n,s\beta}$ sample. So the trace of the summation is small. By the linearity of trace and from Equation (2),

   $$\mathsf{Tr}(\boldsymbol{C}) = p\mathsf{Tr}(\sum_{i \in [s]} \boldsymbol{\phi}(\boldsymbol{c}_i(\boldsymbol{x}))) + \mathsf{Tr}(m)$$

   Since $\mathsf{Tr}(1) = n$ and p is not a divisor of n, $\mathcal{A}$ breaks the semantic game for the encryption $\boldsymbol{C}$ as below.

   $$\begin{aligned} \mathsf{Tr}(\boldsymbol{C}) \bmod p =& 0, \text{ Return } \boldsymbol{C} \text{ is an encryption of } 0 \\ =& n \bmod p = 1, \text{ Return } \boldsymbol{C} \text{ is an encryption of } 1 \end{aligned}$$

[5] Note that the attack does not use the homomorphic property of the encryption scheme, just regular encryptions of bits.

2. When $p = 2$ is a divisor of n.
   In this case, $\mathsf{Tr}(\boldsymbol{C}) \bmod p$ will be 0 for both encryptions of 0 and 1. To get a semantic attack, $\mathcal{A}$ needs to do a small modification here.
   $\mathcal{A}$ picks an $\mathsf{FFI}_{q,1,n,\beta}$ sample $\boldsymbol{C}^*$ (with pre-image $\boldsymbol{c}^*$) such that $|\mathsf{Tr}(\boldsymbol{C}^*)|$ is not a multiple of p. As any field isomorphism map elements in $\mathbb{F}_q$ to itself, this happens with probability $1/p$ for each $p^{-1}\boldsymbol{C}_i(\boldsymbol{y})$, where $\boldsymbol{C}_i(\boldsymbol{y})$ for $1 \leq i \leq S$ are the public key samples, so $\mathcal{A}$ almost surely knows such a sample.
   Multiplying both sides of Equation (2) by $\boldsymbol{C}^*$, we get

$$\boldsymbol{C}\boldsymbol{C}^* = p\ \boldsymbol{\phi}\left((\sum_{i\in[s]} \boldsymbol{c}_i(\boldsymbol{x}))\boldsymbol{c}^*\right) + m\boldsymbol{C}^* \tag{3}$$

   By the ternary sparse choice of the minimal polynomial of $\boldsymbol{x}$, the noise of polynomial multiplication in $\mathbb{X}$ grows at most by a factor of n^3 (Equation 5 of [4]). So, the product of $\boldsymbol{\phi}(\sum_{i\in[s]} \boldsymbol{c}_i(\boldsymbol{x}))$ and $\boldsymbol{\phi}(\boldsymbol{c}^*)$ is an $\mathsf{FFI}_{q,1,n,0.28s\beta^2n^3}$ sample, and hence, by Lemma 2, the absolute value of the trace is bounded by $0.11s\beta^2n^5$. By the linearity of trace, we have from Equation (3)

$$\mathsf{Tr}(\boldsymbol{C}\boldsymbol{C}^*) = p\ \mathsf{Tr}\left(\boldsymbol{\phi}\left((\sum_{i\in[s]} \boldsymbol{c}_i(\boldsymbol{x}))\boldsymbol{c}^*\right)\right) + m\mathsf{Tr}(\boldsymbol{C}^*)$$

   Since $\mathsf{Tr}(\boldsymbol{C}^*)$ is not a multiple of p, $\mathcal{A}$ breaks the semantic game for the encryption $\boldsymbol{C}$ as below.

$$\begin{aligned} \mathsf{Tr}(\boldsymbol{C}\boldsymbol{C}^*) \bmod p =& 0,\ \text{Return } \boldsymbol{C} \text{ is an encryption of 0} \\ =& 1,\ \text{Return } \boldsymbol{C} \text{ is an encryption of 1} \end{aligned}$$

If the q is chosen as in the Theorem to avoid modular reduction, $\mathcal{A}$ returns $m' \in \{0, 1\}$ with advantage $\delta = 1$. The adversary $\mathcal{A}$ runs in polynomial time by the argument given in Theorem 1.

Below we also provide an alternative approach to break the semantic security of the homomorphic encryption scheme that gives a tighter lower bound on q. The main ingredient of this approach is to break the semantic security of the scheme by the distribution of the trace of the inverse of p.

Theorem 3. *Let $q > 0.44s\beta^2n^5 + 0.39\beta n^2$, then there exists a deterministic polynomial time attack against the semantic security of the fully homomorphic encryption scheme $\mathcal{E}$ defined as above.*

Proof. Let the challenger $\mathcal{C}$ give the public key $\mathsf{pk} \hookleftarrow \mathsf{KeyGen}(1^\lambda)$ and an encryption $\boldsymbol{C}$ of a message $m \hookleftarrow \{0, 1\}$ to the adversary $\mathcal{A}$. $\mathcal{A}$ computes

$$\boldsymbol{C}p^{-1} = \sum_{i\in[s]} \boldsymbol{\phi}(\boldsymbol{c}_i(\boldsymbol{x})) + mp^{-1} \tag{4}$$

$\mathcal{A}$ wins the semantic game by the following analysis.

We consider the following two cases for the choice of n.

1. When $p = 2$ is not a divisor of n.
   Since any field isomorphism $\boldsymbol{\phi}$ map elements in $\mathbb{F}_q$ to itself, $\sum_{i\in[s]} \boldsymbol{\phi}(\boldsymbol{c}_i(\boldsymbol{x}))$ is an $\mathsf{FFI}_{q,1,n,s\beta}$ sample. As a consequence, from Equation (4), $\boldsymbol{C}p^{-1}$ is an $\mathsf{FFI}_{q,1,n,s\beta}$ sample for an encryption of 0, thus have small trace. But for an encryption of 1, by linearity, trace of $\boldsymbol{C}p^{-1}$ is dominated by the trace of $1/p$. We now claim the absolute value of $\mathsf{Tr}(1/p)$ is close to the boundary point $(q-1)/2$ of the $\mathbb{F}_q$ representation, thus trace of $\boldsymbol{C}p^{-1}$ is large.
   By the definition of trace,

$$\begin{aligned}\mathsf{Tr}(1/p) &= n/p \bmod q \\ &= n(q+1)/p \bmod q \\ &= (n+q)/p \bmod q \text{ since } p \text{ is not a divisor of } n\end{aligned} \tag{5}$$

   To see the validity of the last line of the above equation, writing $n = pi + 1$, we have

$$\begin{aligned}n(q+1)/p \bmod q &= (pi+1)(q+1)/p \bmod q \\ &= i + (q+1)/p \bmod q \\ &= (n+q)/p \bmod q\end{aligned}$$

   Finally, we have $|\mathsf{Tr}(1/p)| = (q-n)/p$ in the representation of $\mathbb{F}_q$. Thus $\mathcal{A}$ breaks the semantic game for the encryption $\boldsymbol{C}$ from Lemma 2 as below.

$$\begin{aligned}|\mathsf{Tr}(\boldsymbol{C}p^{-1})| \leq & 0.39s\beta n^2, \text{ Return } \boldsymbol{C} \text{ is an encryption of } 0 \\ & \text{Otherwise, Return } \boldsymbol{C} \text{ is an encryption of } 1\end{aligned}$$

2. When $p = 2$ is a divisor of n.
   In this case, from the first line of Equation (5), the trace of $\boldsymbol{C}p^{-1}$ will be small for both encryptions of 0 and 1. To get a semantic attack, $\mathcal{A}$ needs to do a small modification here as in the previous Theorem.
   $\mathcal{A}$ picks an $\mathsf{FFI}_{q,1,n,\beta}$ sample $\boldsymbol{C}^*$ such that $|\mathsf{Tr}(\boldsymbol{C}^*)|$ is not a multiple of p. Multiplying both sides of Equation (4) by $\boldsymbol{C}^*$, we get

$$\boldsymbol{C}^*\boldsymbol{C}p^{-1} = \boldsymbol{C}^* \sum_{i\in[s]} \boldsymbol{\phi}(\boldsymbol{c}_i(\boldsymbol{x})) + m\boldsymbol{C}^*p^{-1} \tag{6}$$

   Again, by the choice of the minimal polynomial of $\boldsymbol{x}$, the noise of polynomial multiplication in $\mathbb{X}$ grows at most by a factor of n^3 (Equation 5 of [4]). For an encryption of 0, since $\boldsymbol{C}^*$ and $\boldsymbol{C}p^{-1}$ are $\mathsf{FFI}_{q,1,n,\beta}$ and $\mathsf{FFI}_{q,1,n,s\beta}$ samples, respectively, $\boldsymbol{C}^*\boldsymbol{C}p^{-1}$ is an $\mathsf{FFI}_{q,1,n,0.28s\beta^2n^3}$ sample. Thus the trace of the product $\boldsymbol{C}^*\boldsymbol{C}p^{-1}$ is still small. But for an encryption of 1, by Equation (6), the trace of $\boldsymbol{C}^*\boldsymbol{C}p^{-1}$ is dominated by the trace of second summand $\boldsymbol{C}^*p^{-1}$. Since the absolute value of $\mathsf{Tr}(\boldsymbol{C}^*p^{-1})$ is close to the boundary point $(q-1)/2$, the trace of $\boldsymbol{C}\boldsymbol{C}^*p^{-1}$ is large in this case.
   To see the above claim, let $\mathsf{Tr}(\boldsymbol{C}^*) = t$, where $|t| \leq 0.39\beta n^2$ and $|t|$ is not a divisor of p. By the linearity of trace and from the previous analysis,

$$\begin{aligned}\mathsf{Tr}(\boldsymbol{C}^*p^{-1}) &= 1/p\mathsf{Tr}(\boldsymbol{C}^*) \\ &= (t+q)/p \bmod q\end{aligned}$$

Finally, we have $|\mathsf{Tr}(\boldsymbol{C}^* p^{-1})| = (q-t)/p$ in the representation of $\mathbb{F}_q$.
By the Equation 5 of [4] and Lemma 2, $\mathcal{A}$ breaks the semantic game for the encryption $\boldsymbol{C}$ as below.

$$|\mathsf{Tr}(\boldsymbol{C}\boldsymbol{C}^* p^{-1})| \leq 0.11 s\beta^2 n^5, \text{ Return } \boldsymbol{C} \text{ is an encryption of } 0$$
$$\text{Otherwise, Return } \boldsymbol{C} \text{ is an encryption of } 1$$

The condition on q in the Theorem ensures $\mathcal{A}$ returns $m' \in \{0,1\}$ with advantage $\delta = 1$. The adversary $\mathcal{A}$ runs in polynomial time by the argument given in Theorem 1.

Effect of the attacks on the recommended parameters: The recommended parameters of [4] (Appendix C, Table 1) for the different level of the (somewhat) fully homomorphic encryption scheme falls within the range of our semantic attacks, except the first level (which has a small q and tolerates very little noise).

6 Proposed Attack on the Computational FFI problem

In this section, we first express the CFFI problem as a lattice problem. The improvement over the previous attack is that here we can avoid the additional combinatorial step (see Subsection 3.1) to solve the CFFI problem. Furthermore, we show how to solve the problem from a small number of shortest lattice vectors. We rely on the dual of the $\boldsymbol{x}$-basis recovered from the shortest vectors of a q-ary lattice generated from FFI samples.

We first define a q-ary lattice for the given FFI samples.

Definition 3 (Trace lattice). *Let $\boldsymbol{A}_1(\boldsymbol{y}), \boldsymbol{A}_2(\boldsymbol{y}), \ldots, \boldsymbol{A}_k(\boldsymbol{y})$ be the $\mathsf{FFI}_{q,k,n,\beta}$ samples for $k > n$. We define a generating matrix $\mathcal{T}$ of order $k \times n$ with coefficients in $\mathbb{F}_q$ and ij-th element defined by $\mathsf{Tr}(\boldsymbol{A}_i(\boldsymbol{y})\boldsymbol{y}^{j-1})$. The q-ary trace lattice is defined as*

$$\mathcal{L}_{\mathcal{T},q} = \{\boldsymbol{\alpha} \in \mathbb{Z}^k : \mathcal{T}\boldsymbol{C} = \boldsymbol{\alpha} \bmod q \text{ for some } \boldsymbol{C} \in \mathbb{Z}^n\}$$

By linearity of trace, the lattice $\mathcal{L}_{\mathcal{T},q}$ contains traces of every finite field element (represented in $\boldsymbol{y}$-basis) with respect to FFI samples $\boldsymbol{A}_i(\boldsymbol{y})$.

Lemma 3. *The q-ary lattice $\mathcal{L}_{\mathcal{T},q}$ has the following properties.*

1. *Its dimension is k.*
2. *Its volume is q^{k-n}.*
3. *It contains n linearly independent vectors $\boldsymbol{\alpha}_i$ such that $\|\boldsymbol{\alpha}_i\| \leq \beta$ for $1 \leq i \leq n$.*

Proof. The first two properties of the Lemma are true for any q-ary lattice of this form. We prove the third point.

For $1 \leq i \leq n$, let $\boldsymbol{C}_{i-1}$ be the dual $\boldsymbol{x}$-basis in the finite field $\mathbb{F}_{q^n}$. Then, recalling the definition of the dual basis,

$$\mathsf{Tr}(\boldsymbol{x}^{j-1}\boldsymbol{C}_{i-1}) = \delta_{i-1}^{j-1}$$

It follows from the linearity of the trace function that any FFI sample $\boldsymbol{A}_j$ has

$$|\mathsf{Tr}(\boldsymbol{A}_j\boldsymbol{C}_{i-1})| \leq \beta$$

Thus the trace lattice contains n linearly independent vectors $\boldsymbol{\alpha}_i$ (corresponding to each $\boldsymbol{C}_{i-1}$), such that

$$\|\boldsymbol{\alpha}_i\| = \|\mathsf{Tr}(\boldsymbol{A}_j\boldsymbol{C}_{i-1})\| \leq \beta,\ 1 \leq j \leq k$$

This concludes the proof.

Since β is reasonably smaller than $(q-1)/2$, the n vectors $\boldsymbol{\alpha}_i$ are very likely the shortest vectors in the lattice $\mathcal{L}_{\mathcal{T},q}$. The lattice vectors $\boldsymbol{\alpha}_i$ have Euclidean norm bounded above by $\beta\sqrt{k}$, which is much smaller than that of the Gaussian heuristics.

Note that the shortest vectors of the trace lattice correspond to the dual $\boldsymbol{x}$-basis, given in the $\boldsymbol{y}$-basis representation. By recomputing the dual of this dual basis, we obtain the $\boldsymbol{x}$-basis in the form of its $\boldsymbol{y}$-basis representation, thus recovering the hidden isomorphism $\boldsymbol{\phi}$. This approach eliminate the $O(n!)$ cost associated with the combinatorial step of the previously mentioned hybrid attack.

In practice, it is generally too costly to find the n shortest vectors in a lattice and thus get the complete $\boldsymbol{C}_i$-basis by just using lattice reduction. When applying BKZ reductions with high block size using aborting techniques [3], which is often seen in cryptanalysis, it is more reasonable to only expect getting a small number of shortest lattice vectors. To account for this, we give a probabilistic approach to recover $\boldsymbol{\phi}$ from a subset of two or more elements of the $\boldsymbol{C}_i$-basis. This is based on the observation that each lattice vector associated with the $\boldsymbol{C}_i$-basis has the same expected norm and is all as likely to appear as a shortest vector while reducing the lattice $\mathcal{L}_{\mathcal{T},q}$. We can thus assume that we are getting random elements from the $\boldsymbol{C}_i$-basis.

Lemma 4. *In a set of $m > 1$ elements, sampled uniformly at random from the set of all the dual $\boldsymbol{x}$-basis, there is, with probability $\Omega(m^2/n)$, at least a pair of dual $\boldsymbol{x}$-basis elements $(\boldsymbol{C}_i, \boldsymbol{C}_j)$ whose quotient gives $\boldsymbol{\phi}$.*

Proof. For the uniform choice of the dual $\boldsymbol{x}$-basis elements, there exists at least a pair of consecutive elements $(\boldsymbol{C}_i, \boldsymbol{C}_j)$ with probability $O(m^2/n)$, i.e a pair with $j = i + 1$.

In the good case for us, this pair with $j = i + 1$ is going to satisfy $\boldsymbol{C}_i = \boldsymbol{x}\boldsymbol{C}_j$ which allows us to compute $\boldsymbol{x}$ as $\boldsymbol{C}_i/\boldsymbol{C}_j$.

To see why, recall that by definition, $\boldsymbol{C}_i$ is the unique element such that, for all $0 \leq k < n$, $\mathsf{Tr}(\boldsymbol{x}^k\boldsymbol{C}_i) = \delta_i^k$. Similarly, $\boldsymbol{C}_j$ satisfies $\mathsf{Tr}(\boldsymbol{x}^k\boldsymbol{C}_j) = \delta_j^k$.

Rewriting $\mathsf{Tr}(\boldsymbol{x}^{k-1}(\boldsymbol{x}\boldsymbol{C}_j))$ for $\mathsf{Tr}(\boldsymbol{x}^k\boldsymbol{C}_j)$, we can check that $\boldsymbol{x}\boldsymbol{C}_j$ already satisfy all necessary conditions needed for $\boldsymbol{C}_i$, except the final one with $k = n$.

However, we now prove that even this final equation is often satisfied. When, it is, we indeed have $\boldsymbol{C}_i = \boldsymbol{x}\boldsymbol{C}_j$. We can compute the missing condition as follows:

$$\mathsf{Tr}(\boldsymbol{x}^{n-1}(\boldsymbol{x}\boldsymbol{C}_j)) = \mathsf{Tr}(\boldsymbol{x}^n\boldsymbol{C}_j) = \mathsf{Tr}(-\boldsymbol{g}(\boldsymbol{x})\boldsymbol{C}_j),$$

the last equality holds because $\boldsymbol{x}^n = -\boldsymbol{g}(\boldsymbol{x}) \pmod{\boldsymbol{f}(\boldsymbol{x})}$. Since $\boldsymbol{g}$ has degree $< n$, $\mathsf{Tr}(-\boldsymbol{g}(\boldsymbol{x})\boldsymbol{C}_j)$ is simply the coefficient of $\boldsymbol{x}^j$ in $-\boldsymbol{g}$. When $\boldsymbol{g}$ is chosen as in Section 3, as a uniform ternary polynomial of degree $\leq \lfloor n/2 \rfloor$, this coefficient is always zero when $j > \lfloor n/2 \rfloor$ and it is zero with probability at least $1/3$ otherwise.

6.1 Lattice Reduction on Trace lattice

In this section, we discuss the experimental results of lattice reduction to find the shortest non-zero vectors in the trace lattice. The parameter $(n, q) = (256, 32771)$ appeared in the level 1 fully homomorphic encryption scheme of [4]. We consider (n, q) close to it for our experiments.

The sample size k, which is the lattice dimension, is the parameter that dominates the running time of lattice reduction. If k is too small, for instance, too close to n, the lattice reduction technique could not extract any meaningful vectors. If k is too large, then the running time of the lattice reduction algorithm is too slow. In our experiments, we choose $k = 2n$.

For small n, it is convenient to recover the shortest vectors with a small block size BKZ algorithm, as expected[6]. But as n increases, the larger block size makes the attack inadequate. To circumvent this, we reduce the lattice dimension by applying a pre-processing lattice reduction step with a smaller block size, whose cost is negligible in the context of the attack.

The choice of the parameters (n, q) in our experiments allow to find "somewhat" short vectors in the trace lattice during the pre-processing step. These short vectors do not correspond to the dual $\boldsymbol{x}$-basis, as expected, but have meaningful properties.

Let $\bar{\boldsymbol{C}}_i$ be the recovered finite field basis corresponding to the short lattice vectors. Heuristically, the $\bar{\boldsymbol{C}}_i$-basis act as a *pseudo* dual $\boldsymbol{x}$-basis in $\mathbb{F}_{q^n}$, i.e.

$$|\mathsf{Tr}(\bar{\boldsymbol{C}}_i\boldsymbol{x}^j)| \lessapprox \beta$$

As a result, for any FFI sample $\boldsymbol{A}_j$

$$|\mathsf{Tr}(\bar{\boldsymbol{C}}_i\boldsymbol{A}_j)| \lessapprox n\beta^2 \tag{7}$$

The recovered $\bar{\boldsymbol{C}}_i$-basis contains information about $\boldsymbol{x}$. To exploit this additional information of $\boldsymbol{x}$, we generate a new integer trace lattice $\mathcal{L}_{\bar{\mathcal{T}}} \subset \mathbb{Z}^k$ of dimension n from the lattice basis $\bar{\mathcal{T}}$ computed using the $\bar{\boldsymbol{C}}_i$-basis (instead of $\boldsymbol{y}$-basis in

[6] For example, we recover the complete dual $\boldsymbol{x}$-basis for the parameter ($n = 100$, $q = 10007$) using BKZ block size 5.

Definition 3). [7] The observation from Equation (7) ensures the basis vectors are unusually small in a (relatively) low dimensional lattice, which allows using stronger lattice reduction algorithms effectively to recover the shortest vectors. The details of our experiments are given in Table 1.

n	q	Pre-Processing	$\bar{\boldsymbol{C}}_i$-basis	Final	Status
200	32771	BKZ 12	✓	BKZ 60	Solved
240	32771	BKZ 20	✓	–	Unsolved
256	32771	BKZ 21	✓	–	Unsolved

Table 1. Experimental results

In general, the running time of a BKZ lattice reduction algorithm is exponential on the blocksize. The authors of [3] successfully perform high blocksize BKZ reductions (with extreme pruning originated in [6]) on different lattices for a small number of rounds under the heuristics that most of the progress of BKZ algorithm is made in the early rounds. In the experiments, we also use a similar approach.

We apply a high block size BKZ algorithm (with extreme pruning) on the lattice basis $\bar{\mathcal{T}}$, aborting regularly to check if some shortest lattice vectors are achieved, continuing otherwise. For $(n = 200, q = 32771)$, we could able to find five shortest vectors within 7 days of running BKZ 60 (aborting regularly) in an Intel Xeon CPU E5-2683 v4 @ 2.10GHz with 1200 MHz processor. The Gram-Schmidt norms of the reduced bases are given in Figure 1. For other parameters, we couldn't find the shortest vectors running the (high block size) aborted BKZ reduction in the fplll software for a couple of months. The application of a more sophisticated lattice reduction approach is beyond the scope of the current paper. We, therefore, invite the cryptanalytic efforts on the other set of parameters, possibly using more advanced lattice reduction tools (example, G6K [1]).

7 Conclusion

In this paper, we illustrate on the FFI problem that having a lattice-reduction approach that fails to solve a problem does not necessarily imply that the problem itself is difficult. Indeed, lattice reduction might not be the optimal strategy to approach it.

[7] It is to be noted that we can always generate arbitrary many samples by doing simple arithmetic from the given FFI samples.

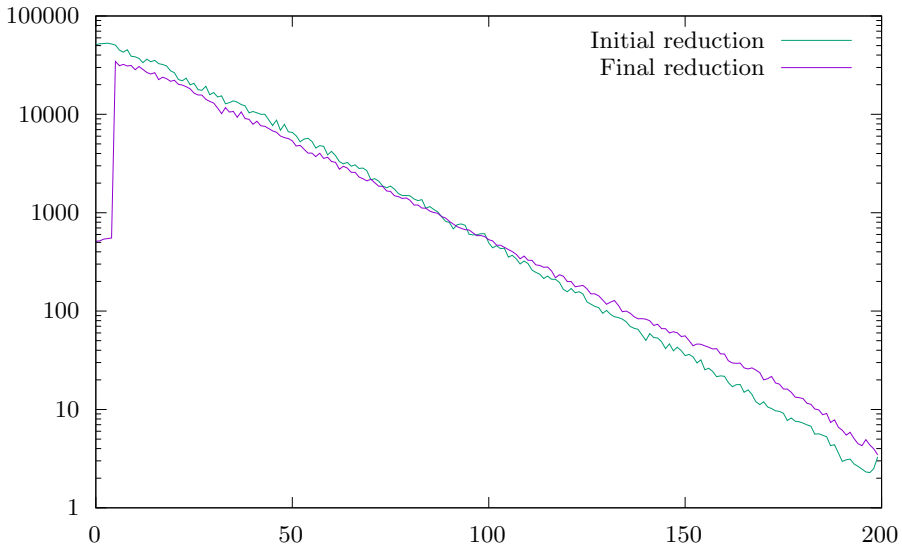

Fig. 1. Gram-Schmidt norms in log scale

Acknowledgements

We thank Anand Kumar Narayanan for helpful discussions on the finite field computation. This work has been supported by the European Union's H2020 Programme under grant agreement number ERC-669891.

References

1. Albrecht, M.R., Ducas, L., Herold, G., Kirshanova, E., Postlethwaite, E.W., Stevens, M.: The general sieve kernel and new records in lattice reduction. In: Ishai, Y., Rijmen, V. (eds.) Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part II. Lecture Notes in Computer Science, vol. 11477, pp. 717–746. Springer (2019). https://doi.org/10.1007/978-3-030-17656-3_25
2. Bosma, W., Cannon, J.J., Steel, A.K.: Lattices of compatibly embedded finite fields. J. Symb. Comput. **24**(3/4), 351–369 (1997). https://doi.org/10.1006/jsco.1997.0138
3. Chen, Y., Nguyen, P.Q.: BKZ 2.0: Better lattice security estimates. In: Lee, D.H., Wang, X. (eds.) Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings. Lecture Notes in Computer Science, vol. 7073, pp. 1–20. Springer (2011). https://doi.org/10.1007/978-3-642-25385-0_1

4. Doröz, Y., Hoffstein, J., Pipher, J., Silverman, J.H., Sunar, B., Whyte, W., Zhang, Z.: Fully homomorphic encryption from the finite field isomorphism problem. In: Abdalla, M., Dahab, R. (eds.) Public-Key Cryptography - PKC 2018 - 21st IACR International Conference on Practice and Theory of Public-Key Cryptography, Rio de Janeiro, Brazil, March 25-29, 2018, Proceedings, Part I. Lecture Notes in Computer Science, vol. 10769, pp. 125–155. Springer (2018). https://doi.org/10.1007/978-3-319-76578-5_5
5. Gama, N., Nguyen, P.Q.: Predicting lattice reduction. In: Smart, N.P. (ed.) Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings. Lecture Notes in Computer Science, vol. 4965, pp. 31–51. Springer (2008). https://doi.org/10.1007/978-3-540-78967-3_3
6. Gama, N., Nguyen, P.Q., Regev, O.: Lattice enumeration using extreme pruning. In: Gilbert, H. (ed.) Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6110, pp. 257–278. Springer (2010). https://doi.org/10.1007/978-3-642-13190-5_13
7. Goldreich, O., Goldwasser, S., Halevi, S.: Public-key cryptosystems from lattice reduction problems. In: Jr., B.S.K. (ed.) Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings. Lecture Notes in Computer Science, vol. 1294, pp. 112–131. Springer (1997). https://doi.org/10.1007/BFb0052231
8. Hoffstein, J., Pipher, J., Silverman, J.H.: NTRU: A ring-based public key cryptosystem. In: Buhler, J. (ed.) Algorithmic Number Theory, Third International Symposium, ANTS-III, Portland, Oregon, USA, June 21-25, 1998, Proceedings. Lecture Notes in Computer Science, vol. 1423, pp. 267–288. Springer (1998). https://doi.org/10.1007/BFb0054868
9. Hoffstein, J., Silverman, J.H., Whyte, W., Zhang, Z.: A signature scheme from the finite field isomorphism problem. J. Math. Cryptol. **14**(1), 39–54 (2020). https://doi.org/10.1515/jmc-2015-0050
10. Kedlaya, K.S., Umans, C.: Fast modular composition in any characteristic. In: 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA. pp. 146–155. IEEE Computer Society (2008). https://doi.org/10.1109/FOCS.2008.13
11. Lenstra, A.K., Lenstra, H.W., Lovász, L.: Factoring polynomials with rational coefficients. Mathematische annalen **261**(ARTICLE), 515–534 (1982). https://doi.org/10.1007/BF01457454
12. Narayanan, A.K.: Fast computation of isomorphisms between finite fields using elliptic curves. In: Budaghyan, L., Rodríguez-Henríquez, F. (eds.) Arithmetic of Finite Fields - 7th International Workshop, WAIFI 2018, Bergen, Norway, June 14-16, 2018, Revised Selected Papers. Lecture Notes in Computer Science, vol. 11321, pp. 74–91. Springer (2018). https://doi.org/10.1007/978-3-030-05153-2_4
13. The FPLLL Development Team: fplll, a lattice reduction library, Version: 5.4.2 (2022), available at https://github.com/fplll/fplll